Lambda Cloud Newsletter: June 2026

Expéditeur : Lambda, Inc. (hello@lambda.ai)
À : k.alahassa123@yahoo.fr
Date : jeudi 2 juillet 2026 à 19:25 UTC

Lambda

Monthly newsletter

This month: Lambda posted the fastest LLM and MoE training times in MLPerf Training v6.0.

We unboxed one of NVIDIA's first co-packaged optics switches and launched Workspaces with scoped access and team-level permissions.

Stephen Balaban joined the MAD Podcast to discuss the real cost stack, from photons to tokens. And Kodiak doubled AI training speed on Lambda's NVIDIA HGX H100 clusters.

Spotlight /

Unbox one of NVIDIA's first co-packaged optics switches with us

Lambda is first to adopt NVIDIA Quantum-X InfiniBand Photonics Q3450-LD switches, bringing co-packaged optics to large-scale AI clusters to reduce network power overhead, lower failure points, and deliver more tokens per watt.

Read more

The GPU myth: state of AI compute 2026

Lambda co-founder and CTO Stephen Balaban joins Matt Turck on the MAD Podcast on the state of AI compute in 2026: what the H100 price indices get wrong, the real cost stack from photons to tokens, and why we're still underbuilding.

Watch the video

Hot off the command line /

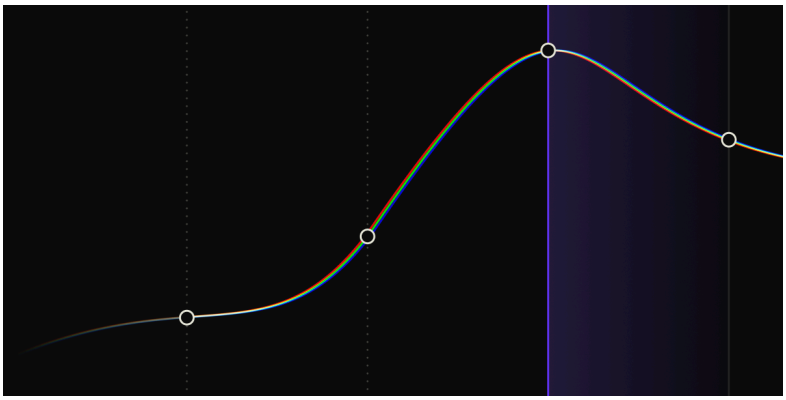

Lambda delivers fastest LLM training on NVIDIA GB300 NVL72 and fastest MoE training on NVIDIA HGX B200

MLCommons' MLPerf Training v6.0 results are in. Lambda posted the fastest pre-training times for both LLM and MoE models. The stack is production-ready. The results are reproducible.

Read more

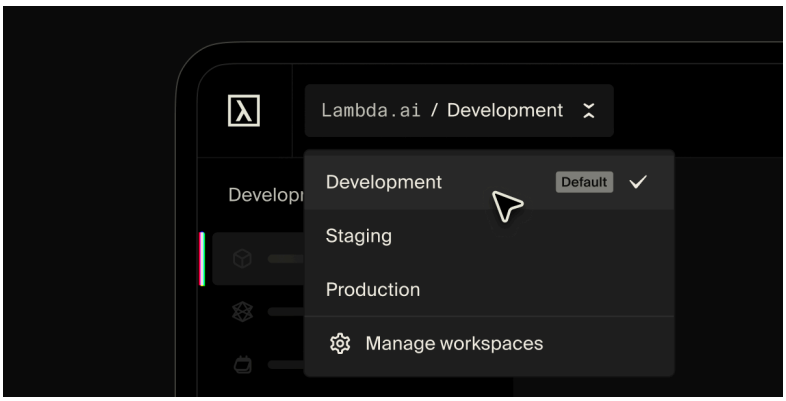

Introducing workspaces for Lambda Cloud

"Trust everyone" isn't an access policy. Workspaces for Lambda Cloud group instances, storage, and users by project. Scope what your team can see and launch. Available now on every account.

Read more

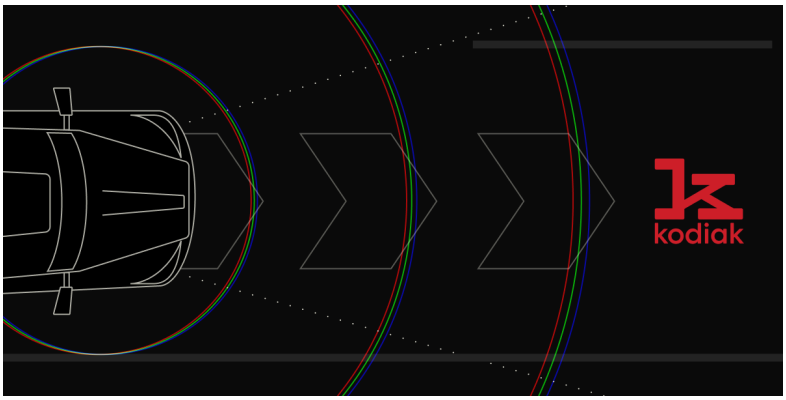

How Kodiak doubled AI training speed on Lambda

28 driverless trucks. The model behind them, GigaFusionNet, trains on Lambda's NVIDIA HGX H100 clusters. 2x experiments at 2x speed.

Watch the video

Knowledge drop /

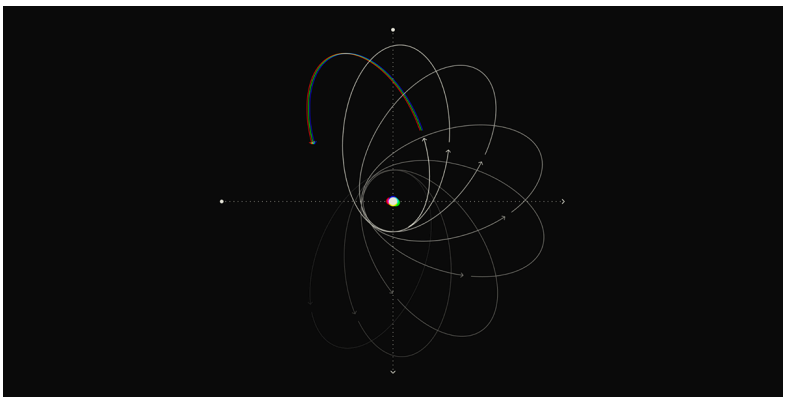

We’re entering the age of large-scale synthetic data

Training data is running out. Human annotation doesn't scale. Sim2Reason generates synthetic physics data that outperforms real-world corpora. +17.9% on JEEBench. Zero human-annotated data. Accepted at ICML 2026. Built by Carnegie Mellon and Lambda.

Read more

3D scene understanding without supervision

3D-Belief learns spatial representations from raw observations, no labels required. A collaboration between Hello Robot, Johns Hopkins, and Cambridge, built on Lambda infrastructure.

Read more

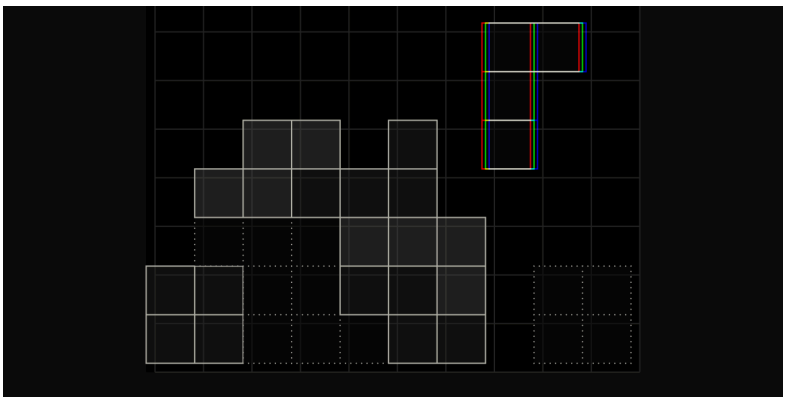

The Lab: a sandbox for agentic experimentation

David Hartmann built The Lab to test how AI agents behave when given real tools on real infrastructure. What Claude Code actually did, what broke, and what the logs revealed.

Read more

Event.log() /

Building at the speed of research: Lambda at CVPR 2026

Lambda at CVPR 2026: two accepted papers, two workshops, and a Kodiak autonomous truck demo. The dominant theme: physical AI. The biggest shift from last year: workloads that combine training, rendering, physics simulation, and inference, not training alone.

Read more

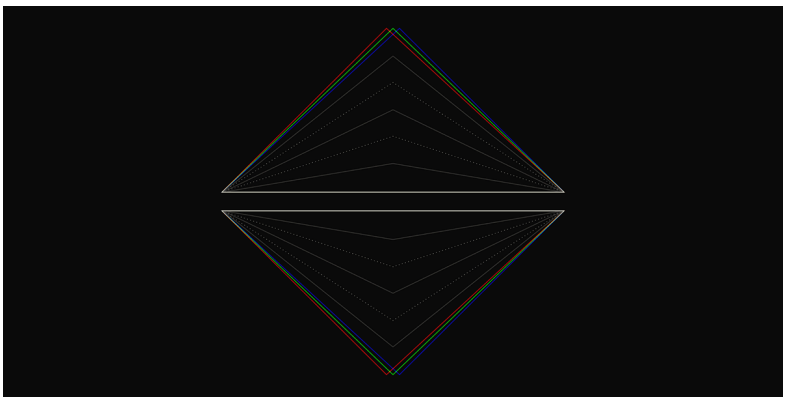

Lambda's keynote at ALVR, co-located with ACL 2026 in San Diego

This week at ACL2026 in San Diego, Lambda's research team is presenting the state of research at the intersection of language, vision, and physical AI.

Read more

Lambda Inc., 2510 Zanker Road, San Jose, California 95131, United States
Unsubscribe | Manage preferences